arXiv:1601.01736v1 [cs.DC] 8 Jan 2016

Algebraic File Synchronization: Adequacy and Completeness

Elod Pal Csirmaz
elod@epcsirmaz.com

Abstract

With distributed computing and mobile applications, synchronizing diverging replicas of data structures is a more and more common problem. We use algebraic methods to reason about filesystem operations, and introduce a simplified definition of conflicting updates to filesystems. We also define algorithms for update detection and reconciliation and present rigorous proofs that they not only work as intended, but also cannot be improved on.

To achieve this, we introduce a novel, symmetric set of filesystem commands with higher information content, which removes edge cases and increases the predictive powers of our algebraic model. We also present a number of generally useful classes and properties of sequences of commands.

These results are often intuitive, but providing exact proofs for them is far from trivial. They contribute to our understanding of this special type of algebraic model, and toward building more complete algebras of filesystem trees and extending algebraic approaches to other data storage protocols. They also form a theoretical basis for specifying and guaranteeing the error-free operation of applications that implement an algebraic approach to synchronization.

Keywords: file synchronization, algebraic approach, distributed systems, reconciliation, data replication

MSC classes: 08A70, 68M07, 68M14, 68P05, 68P20

ACM classes: D.4.3, E.5, F.2.2, G.2

1 Introduction

Synchronization of data structures is a mechanism behind many services we take for granted today: accessing and editing calendar events, documents and spreadsheets on multiple devices, version control systems and geographically distributed internet and web services that guarantee a fast and responsive user experience.

In this paper we investigate a command-based approach to synchronizing filesystem trees. Presented with multiple copies (replicas) of the same filesystem that have been independently modified, we regard the aim of the synchronizer to modify these replicas further to make them as similar as possible. The synchronization algorithm we describe follows the two main steps described by Balasubramaniam and Pierce [1]: update detection, where we identify modifications that have been applied to the replicas since they diverged and represent them as sequences of commands; and reconciliation, where we identify commands that can be propagated to other replicas and do so. The remaining commands, which could not be propagated, represent conflicting updates; resolving these requires a separate system or human intervention.

In this paper we build on our previous work done with Prof. Norman Ramsey on an algebraic approach to synchronization [11], and add to the theoretical understanding of synchronization by providing rigorous proofs that the algorithms we describe work as intended and offer a solution in our framework that cannot be improved on.

A central problem of command-based synchronizers during both update detection and reconciliation is the ordering (scheduling) of commands. If update detection is based on comparing the original state of the filesystem to its new state, then we can easily collect a set of updates, but we need to order these in a way that they could be applied to the filesystem without causing an error (for example, a directory needs to be created before a file can be created under it). Similarly, during reconciliation when we consider all the commands that have been applied to the different replicas, we need to find a way of merging these and creating a global ordering that we can apply to the filesystems. In fact, an insight in [11] is that if the ordering changes the effect of two commands, then under certain circumstances they are not compatible and will give rise to conflicts. Accordingly, in this paper the commutativity of command pairs and their other properties play a special role.

In terms of the usual classification of synchronizers [13, 9, 12], our approach to reconciliation is therefore operation-based as we investigate sets of commands, while the update detection algorithm makes the full synchronizer similar to state-based synchronizers as its starting point is the states of the replicas themselves. Moreover, our ordering algorithms only depend on properties of commands such as commutativity or idempotence, and therefore they can be classified as semantic schedulers. This is in opposition to syntactic scheduling, which uses the origin and the original temporal order of updates, and which may give rise to otherwise resolvable conflicts [12]. In general, ordering plays a central role in operation-based synchronizers. We refer to an excellent survey by Saito and Shapiro of so-called optimistic replication algorithms [12]; and, as examples, to IceCube, where multiple orders are tested to find an acceptable one ([6], see also [7]), or Bayou, where reconciling updates happens by redoing them in a globally determined order [14].

We start by defining a model of filesystems and a set of commands that we will use to describe updates and modifications in Section 2. This is followed by investigating the properties and behaviors of command pairs in Section 3. Section 4 describes update detection, and it is in Section 4.1 that we describe how commands can be ordered and prove that under some simple conditions all valid orders have the same effect.

Section 5 defines our reconciliation and conflict detection algorithms. We then proceed to prove that the output of reconciliation is correct in the sense that it can be applied to the replicas without causing an error, and it is maximal inasmuch as no further updates can be applied under any circumstances. In order to be able to do this, we introduce the concept of *conditional operation* of sequences of commands, and show that it has a number of highly convenient properties. We finish by comparing our results to other research and outlining directions for further work on the introduced algebraic system.

The approach in the present paper offers an improvement over results presented in [11] in multiple ways. We introduce a new set of commands that is symmetric and captures more information about the updates. This has the effect that reasoning about commands becomes simpler as there are fewer edge cases, and more powerful as predictions made by the reasoning are more accurate due to the additional information content. In fact, the new command set not only simplifies our reconciliation algorithm, but also makes it maximal.

The previous work also lacked proofs that the update detection and reconciliation algorithms it presented work as intended, which we provide in the current paper. While the results are intuitive, providing rigorous proofs is far from trivial. During the process, we define a number of useful concepts and show their relationships and the properties they possess. In our view these construct a special algebraic model that is worthy of interest and of further research on its own.

2 Definitions

2.1 The Filesystem

We model a working filesystem using a function Φ with a set of nodes (filesystem paths) $\mathbb{N}$ as its domain, and a set of possible contents or values $\mathbb{V}$ as its codomain. Alternatively, if the filesystem has encountered a problem and is no longer usable, we assign it the special value BROKEN:

Definition 1 (Filesystem)**.**

$$\Phi = \begin{cases} \mathbb{N} \to \mathbb{V} \\ \text{BROKEN} \end{cases}$$

In our model, $\mathbb{N}$ serves as a namespace or "skeleton" for the filesystem. It contains all possible nodes, including the ones where the file system contains no file or directory, where in our model Φ has a special value, $\perp \in \mathbb{V}$.

The nodes form a disjoint union of rooted directed trees, which determines the following PARENT function, and the *ancestor / descendant* and *incomparable* relations. Tao et al. [13] describe a similar filesystem model, although they also model inodes and restrict the filesystem to just a single tree.

Definition 2 (PARENT)**.** The function PARENT(n) returns the parent node of n, or returns NONE if n is the root of a tree.

Definition 3 ($n \prec m$)**.** On $\mathbb{N}$ the *ancestor / descendant* relation is the partial ordering determined by the PARENT function. We write $n \prec m$ if n is the ancestor of m, that is, $n = \text{PARENT}^n(m)$ for some integer $n \geq 1$. We write $n \preceq m$ if $n \prec m$ or $n = m$.

Definition 4 ($n \between m$)**.** We write $n \between m$, or n and m are *incomparable* if $n \not\preceq m$ and $n \not\succeq m$; that is, incomparable nodes are on different branches or on different trees.

Every working filesystem has a so-called *tree property*, which means that if the filesystem is not empty at a node, and the node has a parent, then there must be a directory at the parent node. To model this, in $\mathbb{V}$ we select another special value to represent directories, as we assume that apart from their location in the filesystem, all directories are equal. We do this because, as Bill Zissimopoulos pointed out [15], we often do not want to consider metadata stored in directories (e.g. permission settings) during synchronization, as these are not generally understood well by users, and, if needed, conflict resolution on these settings can be easily automated. We will also see that this assumption makes it possible to define a maximal reconciliation algorithm.

Let therefore the value representing directories be d, and let us define subsets of $\mathbb{V}$ for each type of value: let $\mathbb{V}_{\mathtt{D}}$ be $\{d\}$, let $\mathbb{V}_{\perp}$ be $\{\perp\}$, and let $\mathbb{V}_{\mathtt{F}}$ contain all remaining values that represent different file contents together with any metadata associated with them:

$$\mathbb{V} = \mathbb{V}_{\perp} \cup \mathbb{V}_{\mathtt{F}} \cup \mathbb{V}_{\mathtt{D}}.$$

The tree property can then be defined as

Definition 5 (Tree property)**.**

$$\begin{aligned} \forall n \in \mathbb{N} : \Phi(n) \neq \perp \\ \wedge\, \textsc{parent}(n) \neq \textsc{none} \\ \Rightarrow \Phi(\textsc{parent}(n)) \in \mathbb{V}_{\mathtt{D}}. \end{aligned}$$

This means that as we move down from the root of a tree of nodes, the types of values we encounter in the filesystem can only change according to the transition diagram in Fig. 1, from directories (`D`) to files (`F`) to empty nodes ($\perp$).

In this paper, Φ and Θ denote filesystems, and n, m and o are nodes in $\mathbb{N}$. We use f for an arbitrary element in $\mathbb{V}_{\mathtt{F}}$, and by v_X we mean a value of type X in $\mathbb{V}_X$.

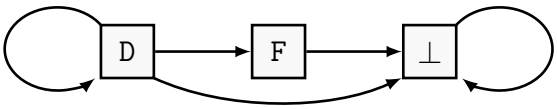

Figure 1: Transitions between types of values in a filesystem

2.2 Commands on Filesystems

Next we define commands on filesystems. As described above, we aimed to select a set of commands that captures as much information about the operations as possible, and is also symmetric. We start by summarizing our reasons for doing so.

Let us consider what kind of information is usually encoded in filesystem operations. A minimal set of commands, based on the most frequent tools implemented by filesystems, may be $create(n, v)$, which creates a file or directory (v) at the empty node n; $edit(n, v)$, which replaces the earlier value at n with v; and $remove(n)$, which deletes the value at n—where $n \in \mathbb{N}$ and $v \in \mathbb{V}$ (but $v \neq \bot$). Regarding their output, that is, the state of the filesystem at n after applying the command, we know that after *create* or *edit*, $\Phi(n) \neq \bot$, whereas after *remove*, $\Phi(n)$ will be $\bot$. However, from [11] and [2] we know that a useful set of axioms will in some cases need to distinguish between, for example, *edit*s that result in directories ($edit(n, d)$) and ones that result in files ($edit(n, f)$), and treat them as separate commands, as their behaviors are quite different when combined with other commands. Indeed, Bill Zissimopoulos' work demonstrated [15] that extending this distinction to more commands ultimately simplifies the definition of conflicting commands, as our model will then able to predict the behavior of commands more precisely. In other words, encoding the type of the output (`D`, `F` or $\bot$) in the commands is definitely useful. At the same time, there is never any need to consider the exact output *value* of a command, as the success or failure of filesystem commands only depends on the types of values in the filesystem.

Notice, however, that the commands listed above also encode some information about their input, the state of the filesystem before the command is applied. In particular, $create(n, v)$ requires that there are no files or directories at n, while $edit(n, v)$ and $remove(n)$ require the opposite. This creates an arbitrary asymmetry where there is now more information available about their output than about their input. As, based on the above, we expect that

encoding more information in the commands results in a model with greater predictive powers, and in order to resolve this asymmetry, we propose a set of commands that encode the type of the original state of $\Phi(n)$ as well. (Some real-life filesystem commands like *rmdir* do this already.)

Definition 6 (Regular commands)**.** Therefore we have four pieces of information related to each command, and represent them using quadruples that list their input type (X), output type (Y), the node to which they are applied (n), and the value they store at the node (v_Y):

$$\langle X, Y, n, v_Y \rangle,$$

where $v_Y \in \mathbb{V}_Y$.

We use `D`, `F` or $\perp$ to represent the different types, and omit the final value if it is determined by the output type (when it is $\perp$ or d). For example, $\langle \perp, \texttt{F}, n, f \rangle$ will represent $create(n, f)$, and $\langle \texttt{D}, \perp, n \rangle$ will represent $rmdir(n)$.

A command can only succeed if the original value at node n has a type that matches the input type of the command. If this is not the case, or if the resulting filesystem no longer has the tree property, then we say that the command *breaks* the filesystem, and assigns BROKEN to it. Broken filesystems are considered equal, but not equal to any working filesystem.

So that we could reason about sequences of commands that break every filesystem, we introduce, in addition to the regular commands, the command `Break` that simply breaks every filesystem.

We note that a command or a sequence of commands is applied to a filesystem by prefixing the command or sequence to it, for example: $\texttt{Break}\,\Phi$, $\langle \perp, \texttt{D}, n, d \rangle\,\Phi$, or simply $S\,\Phi$ if S is a sequence of commands.

Definition 7 (Effect of commands)**.** We use a *replacement operator* of the form $\Phi[v/n]$ to denote a filesystem derived from Φ by replacing its value at n with v:

$$\Phi[v/n](m) = \begin{cases} v & \text{if } m = n \\ \Phi(m) & \text{otherwise.} \end{cases}$$

The effect of the commands is then as follows:

$$\texttt{Break}\,\Phi = \text{BROKEN}$$

$$\langle X, Y, n, v_Y \rangle\, \Phi = \begin{cases} \text{BROKEN} & \text{if } \Phi = \text{BROKEN} \\ \text{BROKEN} & \text{if } \Phi(n) \notin \mathbb{V}_X \\ \text{BROKEN} & \text{if } \Phi[v_Y/n] \text{ violates the tree property} \\ \Phi[v_Y/n] & \text{otherwise.} \end{cases}$$

In general, when we describe multiple or unknown commands, we may substitute one or both of their types with variables, as we did with $\langle X, Y, n, v_Y \rangle$. In this context, if we write that $\langle X, Y \rangle = \langle Z, W \rangle$, we mean that the input and output types of the commands are the same ($X = Z$ and $Y = W$), while a full equality ($\langle X, Y, n, v_Y \rangle = \langle Z, W, m, v_W \rangle$) implies that their nodes and output values are also the same.

For reasons also listed in [11], in this paper we will not consider a *move* or *rename* command. Regarding the theoretical reasoning we aim to follow, this turns out to be useful because this would be the only command that affects filesystems at two nodes at once, therefore describing the dependencies for *move* would call for a more complicated model. From a pragmatic perspective, this restriction does not mean that in an application implementing conflict resolution using the algorithm described here would not be able to handle renames by pre- and post-processing changes in the filesystem to discover them, which can enhance usability, but which (especially when a rename is combined with changes to the content) is a non-trivial problem in itself.

3 Investigating Command Pairs

So that we can describe the effects of commands in general, let us introduce some notation and note some observations. We already know that commands usually do not occur in isolation, and are applied to filesystems in time. Therefore we will investigate sequences of commands with a well-defined order.

Definition 8 (Sequences of commands: $\circ$ and λ)**.** We use "$\circ$" to concatenate commands to form a sequence, or concatenate sequences to form a longer

sequence: $S = \langle \bot, \mathtt{D}, n, d \rangle \circ \langle \bot, \mathtt{F}, m, f \rangle$, or $S' = S \circ \langle \bot, \bot, o \rangle$, with the meaning that the command on the left is applied first:

$$\alpha \circ \beta \, \Phi = \beta(\alpha \, \Phi).$$

Sequences of commands form a free semigroup, and, as usual, λ denotes its unit element, the empty sequence, which by definition leaves all filesystems unchanged.

The following two relations (which echo the ones defined in [11]) allow us to investigate sequences independently of filesystems. In the definitions, A, B, S, T stand for arbitrary sequences of commands.

Definition 9 ($\equiv$)**.** $A \equiv B$ means that A and B are equivalent, that is, they behave in the same way on all filesystems: $\forall \Phi : A \, \Phi = B \, \Phi$. It is a reflexive, commutative and transitive relation, and clearly $A \equiv B \Rightarrow S \circ A \circ T \equiv S \circ B \circ T$.

Definition 10 ($\sqsubseteq$)**.** $A \sqsubseteq B$ means that B extends A, that is, they behave in the same way on any filesystem A does not break: $\forall \Phi : A\Phi \neq \textsc{broken} \Rightarrow A \, \Phi = B \, \Phi$. It is a reflexive and transitive relation, and we can see that $A \sqsubseteq B \Rightarrow S \circ A \circ T \sqsubseteq S \circ B \circ T$, and $A \equiv B \Rightarrow A \sqsubseteq B$.

Our aim is to derive information about the effects of sequences of commands independently of the actual filesystems. In order to make this possible, we investigate the smallest building blocks of sequences: pairs of commands that act on a filesystem directly one after the other. This approach is useful as there are a limited number of command pairs, because, as we argued above, we can disregard the exact output values of commands apart from their type, and we can also abstract the relationship between the nodes in the two commands to a finite number of cases.

We methodically investigated of all possibilities using a computer program[1] to determine which pairs of commands cause errors all the time, which can be simplified to one or no commands, and which commute and can be reversed without any change in their overall effect. Our basis for this investigation was the model of filesystems introduced in this paper. These properties of command pairs are crucial as they determine how a set of commands can be re-ordered to be applied to a filesystem during synchronization,

[1]The program is accessible on-line at
https://github.com/csirmaz/AlgebraicSyncPaper/blob/master/p2/prove.py.

and what commands will never be compatible. Below we list a number of statements derived using this method.

Pairs of commands in general have the form

$$\langle X, Y, n, v_Y \rangle \circ \langle Z, W, m, v_W \rangle$$

where $X, Y, Z, W \in \{\mathtt{D}, \mathtt{F}, \bot\}$, $n, m \in \mathbb{N}$, and values are of the appropriate type: $v_Y \in \mathbb{V}_Y$ and $v_W \in \mathbb{V}_W$.

Rule 1. Commands on incomparable nodes commute: $\langle X, Y, n, v_Y \rangle \circ \langle Z, W, m, v_W \rangle \equiv \langle Z, W, m, v_W \rangle \circ \langle X, Y, n, v_Y \rangle$ where $n \wr\!\!\wr m$.

Rule 2. Commands on incomparable nodes do not break every filesystem: $\langle X, Y, n, v_Y \rangle \circ \langle Z, W, m, v_W \rangle \not\equiv \mathtt{Break}$ where $n \wr\!\!\wr m$.

Rule 3. Commands on the same node break every filesystem if their types are incompatible: $\langle X, Y, n, v_Y \rangle \circ \langle Z, W, n, v_W \rangle \equiv \mathtt{Break}$ where $Y \neq Z$.

Rule 4. Commands on the same node simplify: $\langle X, Y, n, v_Y \rangle \circ \langle Z, W, n, v_W \rangle \sqsubseteq \lambda$ where $Y = Z$, and $X = W = \bot$ or $X = W = \mathtt{D}$.

Rule 5. Commands on the same node simplify: $\langle X, Y, n, v_Y \rangle \circ \langle Z, W, n, v_W \rangle \equiv \langle X, W, n, v_W \rangle$ where $Y = Z$ and $\langle X, W \rangle \neq \langle \bot, \bot \rangle$ and $\langle X, W \rangle \neq \langle \mathtt{D}, \mathtt{D} \rangle$.

Rule 6. Commands on distant relatives break all filesystems: $\langle X, Y, n^*, v_Y \rangle \circ \langle Z, W, n, v_W \rangle \equiv \mathtt{Break}$ and $\langle Z, W, n, v_W \rangle \circ \langle X, Y, n^*, v_Y \rangle \equiv \mathtt{Break}$ where $n^* \prec n$ and $n^* \neq \textsc{parent}(n)$ and $\langle X, Y \rangle \neq \langle \mathtt{D}, \mathtt{D} \rangle$ and $\langle Z, W \rangle \neq \langle \bot, \bot \rangle$.

Definition 11 (Construction pair)**.** A pair of commands on nodes n^* and n is a construction pair if $n^* = \textsc{parent}(n)$ and the commands are one of the following:

$$\langle \bot, \mathtt{D}, n^*, d \rangle \circ \langle \bot, \mathtt{F}, n, f \rangle$$
$$\langle \bot, \mathtt{D}, n^*, d \rangle \circ \langle \bot, \mathtt{D}, n, d \rangle$$
$$\langle \mathtt{F}, \mathtt{D}, n^*, d \rangle \circ \langle \bot, \mathtt{F}, n, f \rangle$$
$$\langle \mathtt{F}, \mathtt{D}, n^*, d \rangle \circ \langle \bot, \mathtt{D}, n, d \rangle$$

Rule 7. All other commands on a child break every filesystem: $\langle X, Y, n^*, v_Y \rangle \circ \langle Z, W, n, v_W \rangle \equiv \mathtt{Break}$ where $n^* = \textsc{parent}(n)$ and $\langle X, Y \rangle \neq \langle \mathtt{D}, \mathtt{D} \rangle$ and $\langle Z, W \rangle \neq \langle \bot, \bot \rangle$ and the pair is not a construction pair.

Definition 12 (Destruction pair)**.** A pair of commands on nodes n and n^* is a destruction pair if $\textsc{parent}(n) = n^*$ and the commands are one of the following:

$$\langle \mathtt{F}, \bot, n\rangle \circ \langle \mathtt{D}, \bot, n^*\rangle$$
$$\langle \mathtt{F}, \bot, n\rangle \circ \langle \mathtt{D}, \mathtt{F}, n^*, f\rangle$$
$$\langle \mathtt{D}, \bot, n\rangle \circ \langle \mathtt{D}, \bot, n^*\rangle$$
$$\langle \mathtt{D}, \bot, n\rangle \circ \langle \mathtt{D}, \mathtt{F}, n^*, f\rangle$$

Rule 8. All other commands on a parent break every filesystem: $\langle X, Y, n, v_Y\rangle \circ \langle Z, W, n^*, v_W\rangle \equiv \mathtt{Break}$ where $\textsc{parent}(n) = n^*$ and $\langle X, Y\rangle \neq \langle \bot, \bot\rangle$ and $\langle Z, W\rangle \neq \langle \mathtt{D}, \mathtt{D}\rangle$ and the pair is not a destruction pair.

Definition 13 (Assertion command)**.** A command is an assertion command if for every filesystem it either breaks it or leaves it in the same state. In other words, assertion commands have the same input and output type, and there must also be only one possible value for their output. Accordingly, $\langle \bot, \bot, n\rangle$ and $\langle \mathtt{D}, \mathtt{D}, n\rangle$ are the only two assertion commands.

Rule 9. An assertion command can be added on a descendant node:

$$\langle \bot, \bot, n\rangle \circ \langle X, Y, n^*, v_Y\rangle \equiv \langle X, Y, n^*, v_Y\rangle \equiv \langle X, Y, n^*, v_Y\rangle \circ \langle \bot, \bot, n\rangle$$

where $n^* \prec n$ and $\langle X, Y\rangle \neq \langle \mathtt{D}, \mathtt{D}\rangle$. (Note that this and next rule is true because the assertion command is adjacent to the original command.)

Rule 10. An assertion command can be added on an ancestor node:

$$\langle \mathtt{D}, \mathtt{D}, n^*\rangle \circ \langle X, Y, n, v_Y\rangle \equiv \langle X, Y, n, v_Y\rangle \equiv \langle X, Y, n, v_Y\rangle \circ \langle \mathtt{D}, \mathtt{D}, n^*\rangle$$

where $n^* \prec n$ and $\langle X, Y\rangle \neq \langle \bot, \bot\rangle$.

Rule 11. Assertion commands can be removed: $\langle X, Y, n, v_Y\rangle \sqsubseteq \lambda$ where $\langle X, Y\rangle = \langle \bot, \bot\rangle$ or $\langle X, Y\rangle = \langle \mathtt{D}, \mathtt{D}\rangle$.

We also define the concept of two commands being independent:

Definition 14 ($A \wr\!\wr B$: Independent commands, sequences and sets)**.** Two commands $\langle X, Y, n, v_Y\rangle$ and $\langle Z, W, m, v_W\rangle$ are independent, and we write

$\langle X, Y, n, v_Y\rangle \wr\!\wr \langle Z, W, m, v_W\rangle$ if they have the same effect in either order and do not break all filesystems:

$$\langle X, Y, n, v_Y\rangle \circ \langle Z, W, m, v_W\rangle \equiv \langle Z, W, n, v_W\rangle \circ \langle X, Y, n, v_Y\rangle \not\equiv \texttt{Break}.$$

For two sequences or unordered sets of commands A and B we write $A \wr\!\wr B$ if for all α in A and all β in B, $\alpha \wr\!\wr \beta$.

It is intentional that we use the same symbol for independent commands and sequences as for incomparable nodes. As we will see in the next Lemma, these two concepts are closely related.

Lemma 1. *Two different commands that are not assertion commands are independent iff the nodes they change are incomparable:*

$$\langle X, Y, n, v_Y\rangle \wr\!\wr \langle Z, W, m, v_W\rangle \Leftrightarrow n \wr\!\wr m$$

if $\langle X, Y, n, v_Y\rangle \neq \langle Z, W, m, v_W\rangle$ *and none of* $\langle X, Y\rangle$ *and* $\langle X, W\rangle$ *is* $\langle \bot, \bot\rangle$ *or* $\langle \texttt{D}, \texttt{D}\rangle$.

Proof. This is easy to show based on the output of the software that investigates the behavior of command pairs. However, this proposition can also be derived from the Rules already listed in the following way.

We proceed in two steps. First, we note that from Rules 1 and 2 we know that if $n \wr\!\wr m$, then $\langle X, Y, n, v_Y\rangle$ and $\langle Z, W, m, v_W\rangle$ must necessarily be independent, and therefore $\langle X, Y, n, v_Y\rangle \not\wr\!\wr \langle Z, W, m, v_W\rangle \wedge n \wr\!\wr m$ cannot be the case.

Next, we prove that

$$\langle X, Y, n, v_Y\rangle \wr\!\wr \langle Z, W, m, v_W\rangle \wedge n \not\wr\!\wr m$$

is also impossible as it would lead to contradiction. As $n \not\wr\!\wr m$ we know that $n = m$, $n \prec m$ or $m \prec n$.

If $n = m$, then from Rule 3 we know $Y = Z$ and $W = X$ as otherwise the commands, in one order or the other, would break all filesystems and they could not be independent. The commands are not assertion commands, and so from Rule 5 we know that $\langle X, Y, n, v_Y\rangle \circ \langle Z, W, n, v_W\rangle \equiv \langle Z, W, n, v_W\rangle \circ \langle X, Y, n, v_Y\rangle$ can only happen if $Y = W$ and $v_Y = v_W$. Therefore $X = Y = Z = W$ which is only possible if both commands are $\langle \texttt{F}, \texttt{F}\rangle$ and have the same output value. This, however, contradicts our assumption that they are different.

If $n \prec m$ or $m \prec n$, then from Rule 6 we know that if they are not directly related, then $\langle X, Y, n, v_Y \rangle \circ \langle Z, W, m, v_W \rangle$ breaks all filesystems, and they cannot be independent. From the construction and destruction pairs and Rules 7 and 8 we also see that even if they are directly related, either $\langle X, Y, n, v_Y \rangle \circ \langle Z, W, m, v_W \rangle$ or $\langle Z, W, m, v_W \rangle \circ \langle X, Y, n, v_Y \rangle$ breaks all filesystems, so they again cannot be independent. □

4 Update Detection

In a command-based reconciliation solution we assume that we have two sequences of commands A and B that have modified a single filesystem Φ yielding two different replicas Φ_A and Φ_B that we need to reconcile. While it is conceivable that the sequences would be based on a record of all operations that modified the two filesystems, in most filesystem implementations such records are not implemented, and therefore we must construct suitable sequences by comparing Φ_A and Φ_B to their common ancestor, Φ. This is called update detection.

A sequence that transforms Φ into Φ_A must contain at least one command on each node that changes between the two states of the filesystem. Also, it is easy to see that this is also sufficient, as there are commands available for any input and output value. Specifically, if at node n, $\Phi(n) = v_X$ and $\Phi_A(n) = v_Y \neq v_X$, then we add the command $\langle X, Y, n, v_Y \rangle$ to A where $v_X \in \mathbb{V}_X$ and $v_Y \in \mathbb{V}_Y$.

To be able to describe such sequences better, let us introduce the following properties:

Definition 15 (Minimal set or sequence)**.** A sequence or set of commands is minimal if it contains at most one command on each node.

Definition 16 (Simple set or sequence)**.** A sequence or set of commands is simple if it is minimal and it does not contain assertion commands.

This update detector therefore yields a simple set of commands because we only add a single command for each node, and we only add commands that are necessary, that is, there will be no assertion commands in the set.

The next step in generating the sequences is to order the commands collected. As this task is at the heart of reconciliation itself independently of update detection, we discuss it in the next section. Then, in Theorem 1, we

prove that the resulting sequence returned by the update detector actually works without breaking the filesystem.

4.1 Ordering commands

We often encounter the case where we only have a set of commands without a specified order. As we have seen above, this can occur after the first stage of update detection, but, more importantly, it is actually the task of the reconciler to determine whether there is an order, and if yes, what order, in which updates from different replicas can be applied to a filesystem.

In this section we describe an algorithm with which all possible orderings of a simple set of commands can be generated, but before doing so, let us consider the following lemma.

Lemma 2. *Given a set of commands that does not contain assertion commands, and which can be applied to a filesystem in some order without breaking it, and which contains commands on n and n^* where $n^* \prec n$, then the set must also contain a command on each node between n and n^*.*

Proof. Without loss of generality, we can assume that $n^* \neq \text{PARENT}(n)$. We prove that under the given conditions, the set must contain a command on $\text{PARENT}(n)$. Then, by reapplying this result, we know that the set must contain commands on every ancestor of n up to n^*.

Furthermore, we prove this proposition for sequences, not sets, as if all sequences must contain a command on $\text{PARENT}(n)$, then so must a set because otherwise there would be no order in which it can be applied to any filesystem.

Let A be a sequence that satisfies the conditions; we therefore know that it contains a command $\langle X, Y, n, v_Y \rangle$ on n and another command $\langle Z, W, n^*, v_W \rangle$ on n^*. We use an indirect proof and assume that there are no commands on $\text{PARENT}(n)$ in A. Next, we create a new sequence $A' \equiv A$ in which $\langle X, Y, n, v_Y \rangle$ and $\langle Z, W, n^*, v_W \rangle$ are next to each other. If they are already next to each other in A, there is nothing to do. Otherwise, consider the command next to $\langle X, Y, n, v_Y \rangle$ in the direction where $\langle Z, W, n^*, v_W \rangle$ is. Let this command be $\langle Q, R, m, v_R \rangle$. If $\langle Z, W, n^*, v_W \rangle$ is to the right, then A looks like the following:

$$A = \cdots \circ \langle X, Y, n, v_Y \rangle \circ \langle Q, R, m, v_R \rangle \circ \cdots \circ \langle Z, W, n^*, v_W \rangle \circ \cdots$$

If $n \wr m$, then swap $\langle X, Y, n, v_Y \rangle$ and $\langle Q, R, m, v_R \rangle$. Based on Rule 1 we know that the new sequence is equivalent to A. Otherwise, we know $m \neq \text{PARENT}(n)$ as there are no commands on $\text{PARENT}(n)$, and so from Rule 6 we get $A \equiv \texttt{Break}$ which contradicts our assumptions. (Note that A cannot contains assertion commands.) By repeating this step we can therefore convert A into A' where $\langle X, Y, n, v_Y \rangle$ and $\langle Z, W, n^*, v_W \rangle$ are neighboring commands. However, then Rule 6 applies to the sequence and therefore $A \equiv A' \equiv \texttt{Break}$ which is again a contradiction. □

We can now define our algorithm to generate orders of sets of commands.

Let U be a simple set of commands that we know can be applied to at least one filesystem in some order without breaking it. Then we can split U into a maximum number of disjunct subsets (components) so that the nodes of two commands from two different subsets are incomparable. These will simply be the components of the subset of the node forest that contains the nodes affected by U.

From Rule 1 we know that commands from different components can be applied in any order. Therefore once we have ordered the commands inside the components, we generate all permutations of the components themselves, including the ones in which commands from separate components mix and overlap (but where the possible orders determined for commands inside the components are respected).

Next, consider the ordering inside a component.

This is not a problem in the case of components containing a single command only. (Note that a file-to-file command ($\langle \texttt{F}, \texttt{F} \rangle$) can only appear in a component of one.)

In each component that is larger than one, we know that all nodes are related, and from Lemma 2 we also know that it must contain commands on nodes in between the related nodes. Therefore the nodes form a rooted ordered tree, and the pairs of commands connected by the edges in the tree must be construction or destruction pairs. (Otherwise, there would be no ordering of the component that is not equivalent to `Break` and via Rule 1 there would be no ordering of the whole set, either.)

To order the commands in the component, select the command with the topmost node, so that there is no command on its parent node (if there is one). If the command removes content at this node (it is of the form $\langle X, \perp \rangle$), then it must follow all other commands in the component, and all parent–

child pairs in the whole component must be destruction pairs. Otherwise, the command with the topmost node needs to precede the other commands, and the component must be made of construction pairs. These follow from the patterns of construction and destruction pairs, which at no point can mix without causing any ordering to break all filesystems.

Then move the command with the topmost node to the beginning or the end of the sequence we are creating out of the component, and recursively start the whole algorithm again on the commands in the component that remain, potentially splitting them into sub-components. The result of this process is a set of possible orders for each component, which are then combined into a set of possible orders for the whole set of commands.

Definition 17 ($\vec{\wp}$)**.** We use $\vec{\wp}(U)$ to denote the set of sequences that contain all commands in the simple set U and represent all orders generated by this algorithm.

It follows from the algorithm that:

Corollary 1. *If U contains commands on both n and* $\textsc{parent}(n)$*, then there is a $U' \in \vec{\wp}(U)$ where they are next to each other.*

As a specific case, we can regard a simple sequence of commands S as a set and determine its $\vec{\wp}(S)$. Lemma 4 is an important property of simple sequences: namely, that all of their valid orderings are equivalent, and all equivalent sequences are mere reorderings. To prove this, first we prove the following:

Lemma 3. *If two simple sequences A and B are equivalent and do not break all filesystems ($A \equiv B \not\equiv$ `Break`), then they must contain the same commands.*

Proof. We use an inverse proof. Let A and B be simple sequences such that $A \equiv B \not\equiv$ `Break`, and Φ be a filesystem that they do not break.

We also assume that they do not contain the same commands. Without loss of generality, we can assume that A contains $\langle X, Y, n, v_Y \rangle$, and B either contains a different command $\langle Z, W, n, v_W \rangle$ on n, or no command on n at all. As A is simple, we know that $\langle X, Y, n, v_Y \rangle$ is not an assertion command, and therefore $X \neq Y$ and $\Phi(n) \neq v_Y$ as it has type X. If so, then if B has no command on n, then $B\,\Phi(n) = \Phi(n)$ and A and B cannot be equivalent. Otherwise we know that $Z = X$ as B does not break Φ either, and that

$v_W = v_Y$ and therefore $W = Y$ as $B\,\Phi(n)$ must be v_Y and B only has one command on n. However, then $\langle Z, W, n, v_W\rangle = \langle X, Y, n, v_Y\rangle$, which is a contradiction. □

Lemma 4. *For all simple sequences T and $S \not\equiv$* `Break`*,*

$$T \equiv S \iff T \in \vec{\wp}(S).$$

Proof. First, as S only has at most one command on each node, any sequence in $\vec{\wp}(S)$ either breaks a filesystem, or it leaves it in the same state as S.

Second, based on the algorithm we also know that $\vec{\wp}(S)$ contains all possible orderings of S as the only ordering between commands it prescribes are when the commands in the reverse order would break every filesystem.

If, for example, in a component or sub-component made up of construction pairs, the command on the topmost node, $\langle X, Y, n, v_Y\rangle$, was not before all other commands in the component as prescribed by the algorithm, then any resulting sequence would necessarily contain a $\langle Z, W, m, v_W\rangle$ where $n \prec m$, potentially followed by commands on incomparable nodes, and then followed by $\langle X, Y, n, v_Y\rangle$. From Rule 1 we know we can convert this sequence to an equivalent sequence by swapping the commands on incomparable nodes with $\langle Z, W, m, v_W\rangle$, which sequence would therefore have the commands on m and n next to each other. Then from Rules 6 and 8 we know the sequence would break all filesystems as a reversed construction pair is never a valid destruction pair.

We can prove that an order not generated by the algorithm in the case of a component made of destruction pairs breaks all filesystems in a similar way.

Also, from Rule 1 we know that the order between any other command pairs is not fixed in the sense that if one order works, so will the other.

From these we know that all sequences in $\vec{\wp}(S)$ are equivalent, and that any sequence of the same commands outside it necessarily breaks all filesystems. We also know $S \in \vec{\wp}(S)$ as $S \not\equiv$ `Break` and therefore $T \in \vec{\wp}(S) \Rightarrow T \equiv S$.

From Lemma 3 we know that if $T \equiv S$, then T contains the same commands as S, from which we see that $T \equiv S \Rightarrow T \in \vec{\wp}(S)$. □

It follows from this that given a simple set, we can simply assume that it has an order. Accordingly, we will at times treat not only sequences as sets, but simple sets as sequences.

4.2 The Correctness of Update Detection

Based on Lemma 2 we can also show that the update detector does function as intended as:

Theorem 1. *A sequence of commands returned by the update detector when comparing the non-broken* Φ^* *to the original* Φ *will not break* Φ *if it is applied to it.*

Proof. Let $\Phi \neq \textsc{broken}$ be the original filesystem, and $\Phi^* \neq \textsc{broken}$ be the filesystem after the changes we intend to detect. Let us assume that some mechanism recorded all changes that occurred to Φ until it reached the state Φ^*. Our set if commands is sufficient to record any change, as there are commands for every input and output type pair, therefore we know that the changes can be recorded as a sequence of commands S where $S\,\Phi = \Phi^*$. Note that S will not contain assertion commands, as they do not represent an actual change in the filesystem.

Therefore, Lemma 2 applies to S and we know if it contains commands on n^* and n where $n^* \prec n$, then it also contains commands on nodes in between; which is equivalent to saying that if Φ changed at node n^* and n, then it had to change on the nodes in between.

If so, then, similarly to what happens in the ordering algorithm, we can separate S into components of commands on related nodes, as commands from separate components freely commute yielding equivalent sequences. From the arguments above we also know that each component must be made up of only destruction or only construction pairs, but not a mixture of the two. This means that in each component, commands on parents either strictly precede or succeed child nodes (as otherwise $S \equiv \texttt{Break}$ would hold), from which it follows that once the components are separated in a new sequence $S' \equiv S$, commands on the same nodes are next to each other in S'.

Applying Rules 4 and 5 we can convert S' into an $S'' \sqsupseteq S'$ in which multiple commands on the same nodes are simplified into a single command or an empty sequence, and therefore S'' is simple and $S \equiv S' \sqsubseteq S''$. It is easy to see that S'' must contain the same commands as the output of the update detector, U, as both are simple but reflect all changes between Φ and Φ^*. If so, then $S'' \in \vec{\wp}(U)$ and based on Lemma 4 we have $S \equiv S' \sqsubseteq S'' \equiv U$, and as $S\,\Phi = \Phi^* \neq \textsc{broken}$, therefore $U\,\Phi \neq \textsc{broken}$. □

An important consequence of this proof is that any sequence of recorded

updates that were applied to a filesystem can be converted to a simple sequence using the rules listed in Section 3 alone.

5 Reconciliation

If we start with two copies of the filesystem Φ, and two different sequences are applied to the copies to yield $\Phi_A = A\,\Phi$ and $\Phi_B = B\,\Phi$, then our aim is to define sequences of commands $A_{\mathcal{R}|B}$ and $B_{\mathcal{R}|A}$ so that $B_{\mathcal{R}|A}\,\Phi_A$ and $A_{\mathcal{R}|B}\,\Phi_B$ would be as close to each other as possible.

We work based on the assumption that to achieve this, we need to apply to Φ_B the commands that have been applied to Φ_A, and *vice versa*. As some commands may have been applied to both filesystems, our first approximation is $A_{\mathcal{R}|B} = A \setminus B$ and $B_{\mathcal{R}|A} = B \setminus A$ (in some suitable order). This, however, will break both filesystems if there have been incompatible updates in A and B. Our aim is therefore to provide an algorithm that selects the commands $A_{\mathcal{R}|B} \subset A \setminus B$ and $B_{\mathcal{R}|A} \subset B \setminus A$ so that $A_{\mathcal{R}|B}\,\Phi_B \neq \textsc{broken}$ and $B_{\mathcal{R}|A}\,\Phi_A \neq \textsc{broken}$, and show that these are the longest sequences with this property, that is, adding any command left out from $A_{\mathcal{R}|B}$ or $B_{\mathcal{R}|A}$ will break the filesystems.

We assume that A and B are returned by an update detector, and so they are simple sequences. (Otherwise, they can be simplified to become simple sequences; see Theorem 1.) We also assume that Φ_A and Φ_B are not broken. With these assumptions, we can now define our reconciliation algorithm.

Definition 18 (Reconciliation)**.** We determine $A_{\mathcal{R}|B}$—which we apply to Φ_B to reconcile it with Φ_A—to be the largest subset of $A \setminus B$ that is independent of $B \setminus A$, where A and B are simple sequences representing the updates applied to replicas Φ_A and Φ_B:

$$
\begin{aligned}
A_{\mathcal{R}|B} = \{\alpha \mid & \\
& \alpha \in A \setminus B \\
& \wedge \\
& \forall \beta \in B \setminus A : \alpha \between \beta \\
\}.&
\end{aligned}
$$

Whether two commands are independent is easy to determine programmatically either based on Definition 14 and the rules listed in Section 3, or

based on Lemma 1 in at most $O(x^2)$ time, where x is the number of commands. Then, to generate the sequence $A_{\mathcal{R}|B}$ from the resulting set, we can order the commands using the algorithm described in Section 4.1 as any subset of A is also simple.

$B_{\mathcal{R}|A}$ can be obtained in a similar way by reversing A and B in the definition.

Next, we would now like to prove that $A_{\mathcal{R}|B}$ and $B_{\mathcal{R}|A}$ can be applied to the replicas, that is, without loss of generality,

$$A_{\mathcal{R}|B}\,\Phi_B \neq \textsc{broken}.$$

So that we can formalize statements needed to prove this, let us introduce two relations that describe under what conditions sequences of commands work, that is, do not break a filesystem.

5.1 Conditional Operation and Inverse Sequences

Definition 19 ($\mathbf{w}(x)$)**.** For a set of k sequences $\mathbf{w}(A_1, A_2, \ldots, A_k)$ means that $A_1, A_2, \ldots, A_k$ work at the same time, that is,

$$\exists\Phi : A_1\Phi \neq \textsc{broken} \wedge \cdots \wedge A_k\Phi \neq \textsc{broken}.$$

As the sequences form a set, their order is irrelevant. We also know that $\mathbf{w}(A) \Leftrightarrow A \not\equiv \texttt{Break}$.

Definition 20 ($\mathbf{w}(x \mid y)$)**.** For two sets of sequences, $\mathbf{w}(A_1, A_2, \ldots, A_k \mid B_1, B_2, \ldots, B_l)$ means that all of $A_1, \ldots, A_k$ work where all of $B_1, \ldots, B_l$ work, that is,

$$\begin{aligned}\forall\Phi : B_1\Phi \neq \textsc{broken} \wedge \cdots \wedge B_k\Phi \neq \textsc{broken}\\ \Rightarrow\\ A_1\Phi \neq \textsc{broken} \wedge \cdots \wedge A_l\Phi \neq \textsc{broken}.\end{aligned}$$

As above, the order of the sequences in the sets is not relevant.

Definition 21 (Sets of sequences)**.** For brevity, we will use calligraphic letters (e.g. $\mathcal{A}, \mathcal{B}, \mathcal{C}$ and $\mathcal{S}$) to denote sets of sequences, and use $A \circ \mathcal{B}$ to denote the set of sequences $\{A \circ B | B \in \mathcal{B}\}$.

It is easy to see that the following corollaries are true:

Corollary 2. $B \sqsubseteq A \Rightarrow \mathbf{w}(A \mid B)$.

Corollary 3. $\forall A, S : \mathbf{w}(A \mid A \circ S)$*, that is, if a sequence works, its initial segment also works.*

Corollary 4. $\mathbf{w}$ *can be chained: if* $\mathcal{A}' \subset \mathcal{A}$*, then* $\mathbf{w}(\mathcal{C} \mid \mathcal{A}') \wedge \mathbf{w}(\mathcal{A} \mid \mathcal{B}) \Rightarrow \mathbf{w}(\mathcal{A} \cup \mathcal{C} \mid \mathcal{B})$.

We also prove the following lemmas.

Rule 12. The combination of independent commands works wherever the original commands work:

$$\begin{aligned}\forall \Phi : \langle X, Y, n, v_Y \rangle &\between \langle Z, W, m, v_W \rangle \\ &\wedge \langle X, Y, n, v_Y \rangle\, \Phi \neq \textsc{broken} \\ &\wedge \langle Z, W, m, v_W \rangle\, \Phi \neq \textsc{broken} \\ &\Rightarrow \langle X, Y, n, v_Y \rangle \circ \langle Z, W, m, v_W \rangle\, \Phi \neq \textsc{broken}.\end{aligned}$$

Proof. We name this proposition a Rule because to prove it, we must reach back to our filesystem model. We proceed in an indirect way and assume that $\langle X, Y, n, v_Y \rangle \circ \langle Z, W, m, v_W \rangle\, \Phi = \textsc{broken}$. We know $\langle X, Y, n, v_Y \rangle\, \Phi \neq \textsc{broken}$ so it must be applying $\langle Z, W, m, v_W \rangle$ that breaks it. Applying a command can only result in a broken filesystem in three cases. First, if the filesystem was already broken, which cannot be the case here. Second, if the input type does not match the filesystem, but we know $\Phi(m) \in \mathbb{V}_W$ and so $(\langle X, Y, n, v_Y \rangle\, \Phi)(m) \in \mathbb{V}_W$ as based on Lemma 1, $n \neq m$. Third, if the new filesystem violates the tree property. This again cannot be the case because we also know that $n \between m$ and the tree property only depends on the types of the parent and children of m, which therefore cannot be changed by $\langle X, Y, n, v_Y \rangle$. □

This result can be extended to sequences:

Lemma 5. *The combination of independent sequences works wherever the original sequences work:*

$$S \between T \Rightarrow \mathbf{w}(S \circ T \mid S, T).$$

Proof. Assume that there is a filesystem Φ so that $S\,\Phi \neq$ BROKEN and $T\,\Phi \neq$ BROKEN, but $(S \circ T)\Phi =$ BROKEN.

From Lemma 1 and Rule 1 we know that the commands in S and T pairwise commute, and so any sequence that contains the commands from S and T and preserve their original partial order is equivalent to $S \circ T$ on all filesystems.

Let the command that breaks Φ in T when applying $S \circ T$ be t so that $T = T_0 \circ t \circ T_1$. It is still true that $(T_0 \circ t)\Phi \neq$ BROKEN, and by definition $(S \circ T_0)\Phi \neq$ BROKEN, but $(S \circ T_0 \circ t)\Phi =$ BROKEN. Also, from above we know that $S \circ T_0 \equiv T_0 \circ S$ and so $(T_0 \circ S)\Phi \neq$ BROKEN.

If we denote the first command in S with s_1, this means that $(T_0 \circ s_1)\Phi \neq$ BROKEN, which we can combine with $(T_0 \circ t)\Phi \neq$ BROKEN, $t \wr\!\!\!\wr s_1$ and Rule 12 (using $T_0\Phi$ as the reference filesystem) to arrive at $(T_0 \circ s_1 \circ t)\Phi \neq$ BROKEN.

We can repeat this step for s_2, the next command in S, and from $(T_0 \circ s_1 \circ t)\Phi \neq$ BROKEN and $(T_0 \circ s_1 \circ s_2)\Phi \neq$ BROKEN arrive at $(T_0 \circ s_1 \circ s_2 \circ t)\Phi \neq$ BROKEN. This can be repeated until S is exhausted and we get $(T_0 \circ S \circ t)\Phi \neq$ BROKEN, which is a contradiction. □

We also prove the following:

Lemma 6. *If A and B are minimal sequences, $\mathbf{w}(A, B)$, and there are commands on node n in both A ($\langle X, Y, n, v_Y\rangle \in A$) and B ($\langle Z, W, n, v_W\rangle \in B$), then the input types of these commands must match ($X = Z$).*

Proof. This result is similar to Lemma 3 and is easily shown using an indirect proof: if $X \neq Z$, then there is no filesystem that either $\langle X, Y, n, v_Y\rangle$ or $\langle Z, W, n, v_W\rangle$ would not break, and consequently A and B cannot work on the same filesystem. □

We continue by defining inverse commands and sequences which allow us to move parts of sequences between the condition and consequence parts of $\mathbf{w}$.

Definition 22 (Inverse commands and sequences)**.** The inverse of command $\langle X, Y, n, v_Y\rangle$ is $\langle X, Y, n, v_Y\rangle^{-1} = \langle Y, X, n, v_X\rangle$ where v_X is an arbitrary value from $\mathbb{V}_X$. We write S^{-1} for the inverse of sequence S, which consists of the inverses of the commands in S in reverse order.

Definition 23 ($\simeq$)**.** For two filesystems we write $\Phi \simeq \Theta$ iff at all nodes the types of their values are the same, or they are both broken.

From the definitions we can clearly see that

Corollary 5. $\forall S, \Phi : (S^{-1})^{-1}\,\Phi \simeq S\,\Phi$.

Lemma 7. *A common initial segment of a set of sequences can be moved to the other side of* $\mathbf{w}$ *by inverting it:*

$$\mathbf{w}(B \circ \mathcal{A} \mid \mathcal{C}) \Rightarrow \mathbf{w}(\mathcal{A} \mid B^{-1} \circ \mathcal{C})$$
$$\mathbf{w}(\mathcal{A} \mid B \circ \mathcal{C}) \Rightarrow \mathbf{w}(B^{-1} \circ \mathcal{A} \mid \mathcal{C})$$

Proof. This is based on the fact that in our model, unless they break a filesystem, commands leave the type of the filesystem value unchanged or change one type into another, but never merge types. That is, sequences—as functions mapping filesystems to filesystems—are essentially bijections over type-equality ($\simeq$) with the only "sink" being BROKEN:

$$\forall S, \Phi, \Theta : S\,\Phi \neq \text{BROKEN} \Rightarrow \\ (S\,\Phi \not\simeq S\,\Theta \iff \Phi \not\simeq \Theta).$$

For the proposition

$$\mathbf{w}(B \circ \mathcal{A} \mid \mathcal{C}) \Rightarrow \mathbf{w}(\mathcal{A} \mid B^{-1} \circ \mathcal{C})$$

the rest of the proof is illustrated by Fig. 2 where $\mathrm{Dom}(x)$, the domain of x, represents the set of filesystems the sequence x does not break, or neither of the sequences in the set x breaks.

We see that the set of filesystems the sequences in $\mathcal{A}$ do not break intersect with the range of B. As B is a bijection between its domain and range, we can use B^{-1} to map this intersection back onto the domain of B. As $\mathbf{w}(B \circ \mathcal{A} \mid \mathcal{C})$ the domain of $\mathcal{C}$ must be a subset of this projected intersection. If so, then we can use B to map the domain of $\mathcal{C}$, which yields the domain of $B^{-1} \circ \mathcal{C}$. As it is also a part of the domain of A, we get $\mathbf{w}(\mathcal{A} \mid B^{-1} \circ \mathcal{C})$.

The second part of the lemma can be proven in a similar way. □

Lemma 8. *The combination of sequences with a common head and independent tails continues to work under the same conditions:*

$$\mathbf{w}(A \circ B \mid \mathcal{S}) \wedge \mathbf{w}(A \circ C \mid \mathcal{S}) \wedge B \between C \Rightarrow \mathbf{w}(A \circ B \circ C \mid \mathcal{S})$$

Proof. Based on Lemma 7 we know $\mathbf{w}(B \mid A^{-1} \circ \mathcal{S})$ and $\mathbf{w}(C \mid A^{-1} \circ \mathcal{S})$, and from Lemma 5 we know $\mathbf{w}(B \circ C \mid B, C)$. Combining these using Corollary 4 we get $\mathbf{w}(B \circ C \mid A^{-1} \circ \mathcal{S})$. Finally, applying the second line of Lemma 7 yields $\mathbf{w}((A^{-1})^{-1} \circ B \circ C \mid \mathcal{S})$, which proves our lemma as $\forall \Phi : (A^{-1})^{-1}\,\Phi \simeq A\,\Phi$ (Corollary 5). □

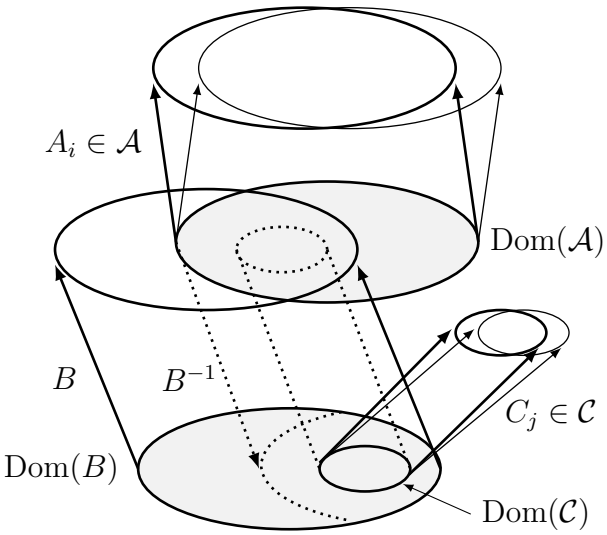

Figure 2: Proof of Lemma 7

5.2 The Correctness of Reconciliation

We are now ready to prove that the proposed algorithm for reconciliation is correct, that is, applying its result is not going to break the replicas. We reformulate the original proposition, $A_{\mathcal{R}|B}\,\Phi_B \neq \textsc{broken}$, based on $\Phi_B = B\,\Phi$ and so we aim to prove that $B \circ A_{\mathcal{R}|B}$ works wherever A and B work:

Theorem 2. *If A, B are simple, then $\mathbf{w}(B \circ A_{\mathcal{R}|B} \mid A, B)$, where, to restate Definition 18,*

$$\begin{aligned} A_{\mathcal{R}|B} = \{\alpha| \\ &\alpha \in A \setminus B \\ &\wedge \\ &\forall \beta \in B \setminus A : \alpha \between \beta \\ \}. \end{aligned}$$

This is trivial unless $\mathbf{w}(A, B)$, so we assume that it is true. Let us first investigate the part of A and B that is excluded from $A_{\mathcal{R}|B}$: their intersection.

Lemma 9. *Let A and B be two simple sequences so that $\mathbf{w}(A, B)$. Then their intersection, in some suitable order, also works on any filesystem either of them works on: $\mathbf{w}(A \cap B \mid A)$, and, by symmetry, $\mathbf{w}(A \cap B \mid B)$.*

Proof. The difficulty is that commands in $A \cap B$ can occur anywhere in A. Therefore, first we prove the following: if all commands are marked in A that also appear in B, then there is an A' in $\vec{\wp}(A)$ in which all marked commands are at the beginning.

We show that A can be transformed via equivalences into A' for which it is true. Our proposition is that if the marked commands are not at the beginning, then the sequence contains an unmarked command followed by a marked command. We show that these can be swapped resulting in an equivalent sequence. Then, by repeating this process similarly to bubble sorting, we can generate a suitable A'.

Let us consider therefore the marked command preceded by an unmarked command in A, and let the marked command be $\langle X, Y, n, v_Y \rangle$, and the preceding unmarked command be $\langle Z, W, m, v_W \rangle$:

$$A = \cdots \circ \langle Z, W, m, v_W \rangle \circ \langle X, Y, n, v_Y \rangle \circ \cdots$$

As A is simple and $\mathbf{w}(A)$, from Rules 6 to 8 we know that these commands can only be on incomparable nodes or form a construction or destruction pair. In the first case, swapping the commands results in a sequence equivalent to A, and we show that the last two cases are impossible as they would lead to contradiction.

In these cases, either $m = \textsc{parent}(n)$ or $n = \textsc{parent}(m)$. If B has a command on m, then from Corollary 1 we know that there is a B' in $\vec{\wp}(B)$ where it is next to $\langle X, Y, n, v_Y \rangle$ (also in B), and from Lemma 4 we know that $B' \equiv B$. When describing our ordering algorithm in Section 4.1 we saw that the command on the parent path determines whether a pair of commands is a construction or destruction pair, and that this, in turn, determines whether the command on the child path must precede or follow the other command as otherwise the sequence would break all filesystems. This argument holds for both A and B, and so the command on m must be on the same side of the command on n in both sequences.

If B has no command on m, then let B' be B, but with an extra command added to it just before $\langle X, Y, n, v_Y \rangle$ according to Rules 9 and 10, from which we know that $B' \equiv B$. We also know that B' is still minimal (but no longer simple).

In either case, therefore, we have a minimal B' which has a command on m just before $\langle X, Y, n, v_Y \rangle$ and $B' \equiv B$. Let this command on m be

$\langle Q, R, m, v_R \rangle$.

$$A = \cdots \circ \langle Z, W, m, v_W \rangle \circ \langle X, Y, n, v_Y \rangle \circ \cdots$$
$$B' = \cdots \circ \langle Q, R, m, v_R \rangle \circ \langle X, Y, n, v_Y \rangle \circ \cdots$$

As $\mathbf{w}(A, B)$ and so $\mathbf{w}(A, B')$, from Lemma 6 we know that $Q = Z$.

Going back to the construction and destruction pairs we see that the output type of the first command is always the same, and it is either $\mathtt{D}$ or $\perp$ depending on the relationship between n and m. Therefore $R = W$ and $\langle Q, R, m, v_R \rangle = \langle Z, W, m, v_W \rangle$ as $|\mathbb{V}_{\mathtt{D}}| = |\mathbb{V}_{\perp}| = 1$. If B originally had a command on m, this is a contradiction as $\langle Z, W, m, v_W \rangle$ was not marked, but we see it must also be in B' and therefore in B. If B had no command on m, this is a contradiction because the command on m in B' is an assertion command, so $Z = Q = R = W$, but A contained no assertion commands.

We now know that there is an $A' \equiv A$ in which commands in $A \cap B$ are at the beginning, and therefore from Corollary 3 we know that $\mathbf{w}(A \cap B \mid A)$ and by symmetry $\mathbf{w}(A \cap B \mid B)$. □

We would like to prove that $\mathbf{w}(B \circ A_{\mathcal{R}|B} \mid A, B)$. From the above we know that we can move commands in $A \cap B$ to the beginning of B and so this is equivalent to

$$\mathbf{w}((A \cap B) \circ (B \setminus A) \circ A_{\mathcal{R}|B} \mid A, B).$$

As we already know that $A_{\mathcal{R}|B} \wr\!\wr (B \setminus A)$ and that $\mathbf{w}((A \cap B) \circ (B \setminus A) \mid A, B)$ (as $\mathbf{w}(B \mid A, B)$), we can prove our theorem based on Lemma 8 if we can prove that

$$\mathbf{w}((A \cap B) \circ A_{\mathcal{R}|B} \mid A, B).$$

Using Lemma 7 and Corollary 5 we can rephrase this as

$$\mathbf{w}(A_{\mathcal{R}|B} \mid (A \cap B)^{-1} \circ A, (A \cap B)^{-1} \circ B).$$

To restate some results above, we therefore already know the following things:

- $(A \setminus B)$ and $(B \setminus A)$ are simple as A and B are simple
- $(A \setminus B) \cap (B \setminus A) = \emptyset$
- $A_{\mathcal{R}|B}$ is the largest subset of $(A \setminus B)$ for which $A_{\mathcal{R}|B} \wr\!\wr (A \setminus B)$

- $\mathbf{w}((A \setminus B), (B \setminus A) \mid (A \cap B)^{-1} \circ A, (A \cap B)^{-1} \circ B)$ because of Lemma 7 as $\mathbf{w}(A, B \mid A, B)$, and $A \equiv (A \cap B) \circ (A \setminus B)$, and $B \equiv (A \cap B) \circ (B \setminus A)$.

Our theorem is therefore proven if we prove the following lemma, in which we rename $A \setminus B$ to S and $B \setminus A$ to T:

Lemma 10. *If*

(c1) S *and* T *are simple sequences,*

(c2) $S \cap T = \emptyset$*,*

(c3) S^* *is the largest subset of* S *where* $S^* \between T$*, and*

(c4) $\mathbf{w}(S, T \mid \mathcal{C})$ *for some set of sequences* $\mathcal{C}$*,*

then

$$\mathbf{w}(S^* \mid \mathcal{C}).$$

Proof. This is trivial unless $\mathbf{w}(\mathcal{C})$, so we will assume that it is true and therefore $\mathbf{w}(S, T)$.

The proof is similar to that of Lemma 9. We mark all commands in S that are in its subset S^*, and we prove that S can be transformed into S' via equivalences so that all marked commands would be at its beginning. If so, then $\mathbf{w}(S^* \mid \mathcal{C})$ based on Corollary 3 as $\mathbf{w}(S' \mid \mathcal{C})$ and $S' = S^* \circ S_0$ where S_0 contains the remaining (conflicting) commands.

Again we know that if S does not already have all marked commands at its beginning, then there is an unmarked command followed by a marked one. We show that these commands are independent and so they can be swapped resulting in an equivalent sequence. By repeating this process we can generate a suitable S'.

Let therefore the marked command in S be $\langle X, Y, n, v_Y \rangle$ and the preceding unmarked command be $\langle Z, W, m, v_W \rangle$. As S is simple and $\mathbf{w}(S)$, from Rules 6 to 8 we know that these commands can only be on incomparable nodes or form a construction or destruction pair. In the first case, swapping the commands results in a sequence equivalent to S, and we show that the other two cases are not possible as they would lead to a contradiction.

In the last two cases, we know that either $m = \textsc{parent}(n)$ or $n = \textsc{parent}(m)$. We also know that because of (c3) there must be a command

$\langle Q, R, o, v_R\rangle$ in T which is not independent of $\langle Z, W, m, v_W\rangle$ as $\langle Z, W, m, v_W\rangle$ is not part of S^*.

$$S = \cdots \circ \langle Z, W, m, v_W\rangle \circ \langle X, Y, n, v_Y\rangle \circ \cdots$$
$$T = \cdots \circ \langle Q, R, o, v_R\rangle \circ \cdots$$

We know none of these commands is an assertion command, and from (c2) that $\langle Q, R, o, v_R\rangle \neq \langle Z, W, m, v_W\rangle$. Therefore, based on Lemma 1 we know this means that either $m \preceq o$ or $o \preceq m$. From (c3) we also know that $\langle X, Y, n, v_Y\rangle \wr\!\!\!\wr \langle Q, R, o, v_R\rangle$, and so because of Lemma 1, $n \wr\!\!\!\wr o$.

We therefore have four cases considering the relationships between n, m and o:

- $n = \textsc{parent}(m) \wedge o \preceq m$. This would mean that $o \preceq n$ or $n = \textsc{parent}(o)$, which contradicts $n \wr\!\!\!\wr o$.
- $n = \textsc{parent}(m) \wedge m \preceq o$. This would mean that $n \prec o$, which contradicts $n \wr\!\!\!\wr o$.
- $m = \textsc{parent}(n) \wedge o \preceq m$. This would mean that $o \prec n$, which contradicts $n \wr\!\!\!\wr o$.
- $m = \textsc{parent}(n) \wedge m \preceq o$. Let us continue with this case.

We know that T cannot have a command on m as $m = \textsc{parent}(n)$ and $\langle X, Y, n, v_Y\rangle \wr\!\!\!\wr T$, and so $m \neq o$ and therefore $m \prec o$. This means we can create a new sequence $T' \equiv T$ by inserting a command on m into T before $\langle Q, R, o, v_R\rangle$ according to Rule 10. We now have $\langle X, Y, n, v_Y\rangle \wr\!\!\!\wr T \equiv T'$, which is a contradiction as T' has a command on m and $m = \textsc{parent}(n)$. □

5.3 Reconciliation is Maximal

The reconciliation algorithm defined above is also maximal, that is, it is not possible to apply any further commands from $A \setminus B$ to Φ_B. To show this, we are going to prove the following theorem:

Theorem 3. *If A and B are simple sequences, and we select a sequence $S \subset (A \setminus B)$ so that it contains a command $\langle X, Y, n, v_Y\rangle$ for which $\langle X, Y, n, v_Y\rangle \not\wr\!\!\!\wr (B \setminus A)$, then $B \circ S$ either breaks all filesystems A and B do not break (that is, S breaks all possible Φ_B replicas), or it changes a node that B has already changed to a different value, that is, it overrides a change in Φ_B.*

Such an override could occur if a given node was modified differently in A and B (to different values but the same type), which is clearly a conflict that should be marked for review by a human or a different system.

Proof. The proof is similar to the ones we have seen above. Without loss of generality, we assume that $\langle X, Y, n, v_Y \rangle$ is the first command in S that is not independent of $B \setminus A$. If so, we can split S into $S_0 \circ \langle X, Y, n, v_Y \rangle \circ S_1$ where $S_0 \wr\!\wr (B \setminus A)$. Let us also split $B \setminus A$ into $T_0 \circ \langle Z, W, m, v_W \rangle \circ T_1$ where $T_1 \wr\!\wr \langle X, Y, n, v_Y \rangle$. From Lemma 9 we know that $B \circ S \equiv (A \cap B) \circ (B \setminus A) \circ S$. We also know that commands in S_0 and $B \setminus A$ commute, and so this is equivalent to $(A \cap B) \circ S_0 \circ (B \setminus A) \circ \langle X, Y, n, v_Y \rangle \circ S_1$. Expanding $B \setminus A$ and swapping T_1 and $\langle X, Y, n, v_Y \rangle$ we get

$$(A \cap B) \circ S_0 \circ T_0 \circ \langle Z, W, m, v_W \rangle \circ \langle X, Y, n, v_Y \rangle \circ T_1 \circ S_1.$$

First, we prove that this sequence breaks all filesystems unless $n = m$. Let us therefore suppose $n \neq m$ and, to use an inverse proof, that the sequence works on Φ where $A\,\Phi \neq \textsc{broken}$. If so, its initial segment, $(A \cap B) \circ S_0 \circ T_0 \circ \langle Z, W, m, v_W \rangle \circ \langle X, Y, n, v_Y \rangle$ must also work.

We know $\langle X, Y, n, v_Y \rangle \not\wr\!\wr \langle Z, W, m, v_W \rangle$, and so from Lemma 1, $n \not\wr\!\wr m$. As the sequence works, $\langle Z, W, m, v_W \rangle \circ \langle X, Y, n, v_Y \rangle$ can only be a construction or destruction pair. As $\langle Z, W, m, v_W \rangle$ and T_0 originate from $B \setminus A$ and B is simple, we also know there are no commands on m in $A \cap B$ or T_0, and that there are no commands on m in S_0 as $S_0 \wr\!\wr (B \setminus A)$.

On the one hand we therefore know that $\Phi(m)$ must be of type Z, but on the other, from Rules 7 and 8, we also know that $\langle X, Y, n, v_Y \rangle$ cannot be applied to Φ without changing $\Phi(m)$ first. As $A\,\Phi \neq \textsc{broken}$, this means that A must contain a command on m. Let this command be $\langle Q, R, m, v_R \rangle$.

Clearly $Q = Z$ (from Lemma 6), and from the construction and destruction pairs we also see that $R = W$ must hold, where $W = \mathtt{D}$ or $W = \bot$. As $|\mathbb{V}_{\mathtt{D}}| = 1$ and $|\mathbb{V}_{\bot}| = 1$, this also means that $v_R = v_W$ and therefore $\langle Q, R, m, v_R \rangle = \langle Z, W, m, v_W \rangle$. This is a contradiction as $\langle Q, R, m, v_R \rangle \in A$, but $\langle Z, W, m, v_W \rangle \in B \setminus A$.

The only possibility is therefore $n = m$. As A and B are simple, we therefore know that these are the first commands on n, and so (again from Lemma 6) $Z = X$. We also know $v_W \neq v_Y$ as otherwise the two commands would be equal and would be included in $A \cap B$, but this means that $\langle X, Y, n, v_Y \rangle$, from A, overrides a change introduced by $\langle Z, W, n, v_W \rangle$, from B, which should be treated as a conflict. □

From the proof it is apparent that this result depends on $|\mathbb{V}_{\mathtt{D}}| = 1$.

6 Related Work

6.1 Liberal and Conservative Reconciliation

We consider the reconciliation algorithm described here an improvement over the one derived during our previous research [11] as the previous algorithm not only fails to propagate all possible commands wherever it is possible (that is, where it does not break the filesystem), but the current algorithm is also simpler. This is because the previous reconciliation algorithm excludes commands from being propagated which must be preceded by a command that conflicts.

The former observation is supported by Berzan and Ramsey, who in [2] describe different general reconciliation policies. The liberal (maximal) policy propagates all updates to all replicas where the update does not break the filesystem, while a conservative policy refrains from updating any node that is below a node with conflicting commands. They show that the reconciliation algorithm in [11] implements an intermediate policy as one can easily construct two- or three-replica scenarios where an update could clearly be propagated, but it is excluded. To rephrase the example in [2], consider the following two update sequences that have been applied to replicas Φ_A and Φ_B:

$$A = \langle \bot, \mathtt{D}, n^*, d_1 \rangle \circ \langle \bot, \mathtt{F}, n, f \rangle$$
$$B = \langle \bot, \mathtt{D}, n^*, d_2 \rangle$$

where $\textsc{parent}(n) = n^*$ ([11] did not require all directories to have the same value). Clearly $\langle \bot, \mathtt{F}, n, f \rangle$ could be applied to Φ_B, but it is not as it must be preceded by creating the directory, which conflicts with the same command in B due to the different value. (Introducing a third replica which has not changed at all further complicates the picture.) The current algorithm no longer needs to specify that there can be no conflicts on preceding commands, and, to use the above terminology, implements a fully liberal policy.

6.2 Comparing Definitions of Conflicts

As a test of the proposed reconciliation algorithm, we compare our definition of conflicting updates to how conflicts are defined by Balasubramaniam and Pierce in their state-based approach implemented in the Unison synchronizer [1]. As we noted in [11], the update detector they describe provides a safe estimate of nodes (paths) at which updates occurred. It marks some nodes as *dirty* in a way that we know that at non-dirty nodes the filesystems (replicas) have not changed between their common original state and current state (see Definition 3.1.1 in [1]). The *dirty* marks are also up-closed, that is, all ancestor nodes of a dirty node are also dirty (Fact 3.1.3 in [1]). And finally, in our notation, a conflict is detected between replicas Φ_A and Φ_B at node n if n is marked as dirty in both Φ_A and Φ_B, and $\Phi_A(n) \neq \Phi_B(n)$, and n does not point to a directory in both replicas. (In other words, $\Phi_A(n) \notin \mathbb{V}_{\mathtt{D}} \vee \Phi_B(n) \notin \mathbb{V}_{\mathtt{D}}$; see section 4.1 in [1]. Let us note that there is an alternative approach to defining an algorithm for the same synchronizer by Pierce and Vouillon in [9].)

It can be easily seen that due to an edge case, not all conflicts detected based on the above definition entails a conflict based on our system; that is, [1] describes a more conservative policy. We use the example we described in [11], where $n^* = \textsc{parent}(n)$ and $\Phi(n^*)$ is a directory and $\Phi(n)$ is a file, and the two replicas are derived in the following way:

$$\Phi_A = \langle \mathtt{F}, \bot, n \rangle \circ \langle \mathtt{D}, \bot, n^* \rangle \, \Phi$$
$$\Phi_B = \langle \mathtt{F}, \bot, n \rangle \, \Phi.$$

Then, n^* is dirty in both replicas: in Φ_A it was modified, and in Φ_B one of its descendants was modified. Moreover, $\Phi_A(n^*) \neq \Phi_B(n^*)$ and $\Phi_A(n^*) = \bot \notin \mathbb{V}_{\mathtt{D}}$, and therefore a conflict is detected at n^*. (This behaviour is preserved in the more recent Harmony synchronizer. See "delete/delete conflicts" in [10, 5].) Our reconciliation algorithm detects no conflicts; instead, it propagates $\langle \mathtt{D}, \bot, n^* \rangle$ to Φ_B, which we think is as expected and desired.

At the same time, it can be shown that if our command-based reconciler detects a conflict, it entails a conflict in the state-based reconciler. We note here that Balasubramaniam and Pierce also suppose that all directories are equal, therefore, as elsewhere, we are safe to continue to assume that $|\mathbb{V}_{\mathtt{D}}| = 1$.

Proof. Let A and B be two simple sequences returned by the update detector for the two replicas Φ_A and Φ_B. A conflict between the commands

$\langle X, Y, n, v_Y\rangle \in A$ and $\langle Z, W, m, v_W\rangle \in B$ means that even though $\langle X, Y, n, v_Y\rangle \in A \setminus B$ and $\langle Z, W, m, v_W\rangle \in B \setminus A$, they cannot be included in $A_{\mathcal{R}|B}$ and $B_{\mathcal{R}|A}$, respectively, because $\langle X, Y, n, v_Y\rangle \not\parallel \langle Z, W, m, v_W\rangle$ (see Definition 18).

From Lemma 1 we therefore know that $n \not\parallel m$. Without loss of generality we can assume that $n \preceq m$. From this we see that n is *dirty* in both Φ_A and Φ_B, as the filesystem changes at n in Φ_A and at n or its descendant in Φ_B (and none of the commands are assertion commands).

Now we only need to show that $\Phi_A(n) \neq \Phi_B(n)$, because from this and $|\mathbb{V}_{\mathtt{D}}| = 1$ we will also know that at least one of these values is not a directory, and therefore there is a conflict at n in the state-based reconciler. If $n = m$, this follows from the fact that $\langle X, Y, n, v_Y\rangle \neq \langle Z, W, m, v_W\rangle$ (because they are not in $A \cap B$), as from Lemma 6 we have $X = Z$, and therefore the above can only be true if $\Phi_A(n) = v_Y \neq v_W = \Phi_B(n)$. If $n \prec m$ and there is no command on n in B, then it is also true as $\langle X, Y, n, v_Y\rangle$ is not an assertion command, and therefore $\Phi_A(n) \neq \Phi(n) = \Phi_B(n)$, where Φ is the common ancestor of the replicas. Finally, if there is a command on n in B, then from $\langle X, Y, n, v_Y\rangle \in A \setminus B$ we know it must be different from $\langle X, Y, n, v_Y\rangle$, and similarly to the first case, we have $\Phi_A(n) \neq \Phi_B(n)$. □

7 Towards a Formal Algebra

Similarly to [11] we can consider creating an algebra over command sequences which will enable us to draw conclusions about the behaviors of the sequences independently of filesystems. This would provide a secondary model of (filesystem) commands above the model a filesystems defined in this paper. The new model would no longer describe filesystems as such, but use known relationships between sequences of commands as its starting point.

In this algebra, equivalence ($\equiv$) and extension ($\sqsubseteq$) would become algebraic relations between sequences, and logical rules involving them (e.g. $A \equiv B \Rightarrow S \circ A \circ T \equiv S \circ B \circ T$) would be re-cast as inference rules. In addition to these, the Rules listed in Section 3 would become axioms that we accept as true, and from which other true statements can be derived using the inference rules. Such a system can allow us to deduce, for example, whether two sequences are equivalent or one extends the other without reference to filesystems.

We note, however, that the Rules to be used as axioms specify the relationships between the nodes the commands act on, which would require a

more complicated set of symbols to represent in the algebra. To avoid this, we think it will be more fruitful to regard commands (and sequences) on different nodes as different, and, in effect, have nine times as many commands in the algebra as there are nodes in the namespace of the filesystems. This allows us to reduce the number of types of symbols in our algebra, and regard the Rules not as axioms, but as templates for axioms from which all axioms (for pairs of separate nodes, etc.) can be generated.

We expect that the soundness and completeness of such an algebra can be proven similarly to the proofs described in [11]. Indeed, results in the current paper can serve as building blocks of such a proof: the computer program used to check the Rules in effect proves the soundness of all axioms, and Lemma 4 proves completeness in a limited sense. It is also worth noting that the majority of proofs presented in this paper do not actually refer to specific filesystems or the definition of commands, but draw on known relationships and the Rules in Section 3. In other words, many proofs are transferable to the algebra that we describe here.

In a more complete algebra from which the correctness and completeness of reconciliation could also be derived, one can also include symbols, inference rules and axioms for $\mathbf{w}(\mathcal{A})$ and $\mathbf{w}(\mathcal{A} \mid \mathcal{B})$. We see proving the soundness and completeness of this extended algebra an intriguing problem that is worthy of further research.

8 Conclusions and Further Research

Apart from constructing an algebra, there are many other ways in which further work can extend the current results. An important extension would be to consider reconciling not only two, but more replicas in a single step and prove the correctness and maximality of the algorithm proposed, or show that it is impossible to satisfy these criteria.

A related problem is to extend the system and the proofs to allow for cases where reconciliation cannot complete fully or if only a subset of the replicas are reconciled (e.g. due to network partitioning), both of which would result in a state where different replicas have different common ancestors, that is, the updates specific to the replicas start from different points in the update history of the filesystem. Existing research can offer pointers as to how such cases can be modelled in our algebraic system. Parker et al. [8] and Cox and Josephson [4] describe version vectors (update histories) kept

as metadata, while Chong and Hamadi present distributed algorithms that allow incremental synchronization [3]. Representing individual updates to files in their modification histories (as described in [4]) as separate commands could also enable an algebraic synchronizer to reconcile otherwise conflicting updates and resolve partial reconciliations.

An extension of our results could also entail relaxing the simplifying assumptions we used. Identifying which results depend on the assumption that $|\mathbb{V}_{\mathtt{D}}| = 1$ and devising updated proofs and algorithms for when $|\mathbb{V}_{\mathtt{D}}| > 1$, or specifying a pre-processing step that reconciles meta-information in directories can enable the synchronizer to identify and resolve a wider range of conflicts. Similarly, designing pre- and post-processors that present removing and creating a node not as separate updates but as a single rename (move), could contribute to the user-friendliness of the reconciler and help avoid human error.

And finally, we hope that this work, together with [11], provides a blueprint of constructing an algebra of commands for different storage protocols (e.g. XML trees, mailbox folders, generic relational databases, etc.), and of demonstrating the adequacy and completeness of update and conflict detection and reconciliation algorithms defined over it. This, in turn, can offer formal verification of the algorithms underlying specific implementations in a variety of synchronizers. Alternatively, by generalizing the parent–child relationships between filesystem nodes, the demonstrated properties of minimal sequences of commands and conditional operation of sequences ($\mathbf{w}$) may also contribute to future research into algebraic structures constrained by predefined sparse connections between their elements.

9 Acknowledgments

I thank Laszlo Csirmaz for his invaluable input on this paper, and Bill Zissimopoulos for his observations on algebraic file synchronization that triggered this research.

References

[1] S. Balasubramaniam and Benjamin C. Pierce, "What is a File Synchronizer?" in *Proceedings of the 4th Annual ACM/IEEE International Con-*

ference on Mobile Computing and Networking. New York, NY: ACM, 1998, pp. 98–108.

[2] Constantin Berzan and Norman Ramsey, "Summer Scholars Technical Report," Tufts University, Medford, MA, 2010. Accessed 29 November 2015. http://thirld.com/files/summerscholars_techreport.pdf.

[3] Yek Loong Chong and Youssef Hamadi, "Distributed Algorithms for Log-based Reconciliation," Microsoft Research, Microsoft Co., Redmond, WA, Tech. Rep. MSR-TR-2004-104, 2004.

[4] Russ Cox and William Josephson, "File synchronization with Vector Time Pairs," MIT CSAIL, Cambridge, MA, Tech. Rep., MIT-CSAIL-TR-2005-014, 2005. Accessed 22 November 2015. http://hdl.handle.net/1721.1/30527.

[5] J. Nathan Foster et al., "Schema-directed data synchronization," U of Pennsylvania, Philadelphia, PA, Technical Reports (CIS), Paper 5, 2005. Accessed 23 November 2015. http://repository.upenn.edu/cis_reports/5.

[6] Anne-Marie Kermarrec and Antony Rowstron and Marc Shapiro and Peter Druschel, "The IceCube Approach to the Reconciliation of Divergent Replicas," in *Proceedings of the Twentieth Annual ACM Symposium on Principles of Distributed Computing.* New York, NY: ACM, 2001, pp. 210–218.

[7] Vidal Martins and Esther Pacitti and Patrick Valduriez, "Distributed semantic reconciliation of replicated data," in *Proc. of CDUR.* Paris, France: CNAM, 2005, pp. 48–53.

[8] D. S. Parker et al., "Detection of Mutual Inconsistency in Distributed Systems," in *IEEE Trans. Softw. Eng.,* vol. 9, no. 3. Piscataway, NJ: IEEE Press, May 1983, pp. 240–247.

[9] Benjamin C Pierce and Jérôme Vouillon, "What's in Unison? A formal specification and reference implementation of a file synchronizer," U of Pennsylvania, Philadelphia, PA, Technical Reports (CIS), Paper 40, 2004. Accessed 23 November 2015. http://repository.upenn.edu/cis_reports/40.

[10] Benjamin C. Pierce, Alan Schmitt, and Michael B. Greenwald, "Bringing Harmony to Optimism: An Experiment in Synchronizing Heterogeneous Tree-Structured Data," U of Pennsylvania, Philadelphia, PA, Technical Reports (CIS), Paper 39, 2004. Accessed 23 November 2015. http://repository.upenn.edu/cis_reports/39.

[11] Norman Ramsey and Elod Csirmaz, "An Algebraic Approach to File Synchronization," in *Proceedings of the Joint 8th European Software Engineering Conference and 9th ACM SIGSOFT Symposium on the Foundations of Software Engineering.* New York, NY: ACM, 2001, pp. 175–185.

[12] Yasushi Saito and Marc Shapiro, "Optimistic Replication," in *ACM Computing Surveys,* vol. 37, no. 1. New York, NY: ACM, March 2005, pp. 42–81.

[13] Vinh Tao and Marc Shapiro and Vianney Rancurel, "Merging Semantics for Conflict Updates in Geo-distributed File Systems," in *Proceedings of the 8th ACM International Systems and Storage Conference.* New York, NY: ACM, 2015, pp. 10:1–10:12.

[14] D. B. Terry et al., "Managing Update Conflicts in Bayou, a Weakly Connected Replicated Storage System," in *Proceedings of the Fifteenth ACM Symposium on Operating Systems Principles.* New York, NY: ACM, 1995, pp. 172–182.

[15] Bill Zissimopoulos, private communication, August 2015.